



Scaling Datacenter Bandwidth with 400G and 800G Ethernet

Addison Chi
Technical Lead
addison@arista.com


Driving a Shift to Cloud Networking



Source: Dell'Oro Market Research, Ethernet Switch Update, October 2016

Google TPU Performance Against CPU/GPU


Google TensorFlow Processor Unit Performance In AI and Machine Learning



CPU, GPU and TPU performance on six reference workloads (in log scale)


Need to Scale Cloud Network Bandwidth 2X/Year

Intra-datacenter Bandwidth Growth




Source: Urs Hoelzle, Google Cloud Jupiter, SIGCOMM 2015

Merchant Network Silicon Growth



Source: The 650 Group, January 2017

Merchant Switch Silicon Bandwidth Growth



Ethernet Speed Transitions by Lane Speed

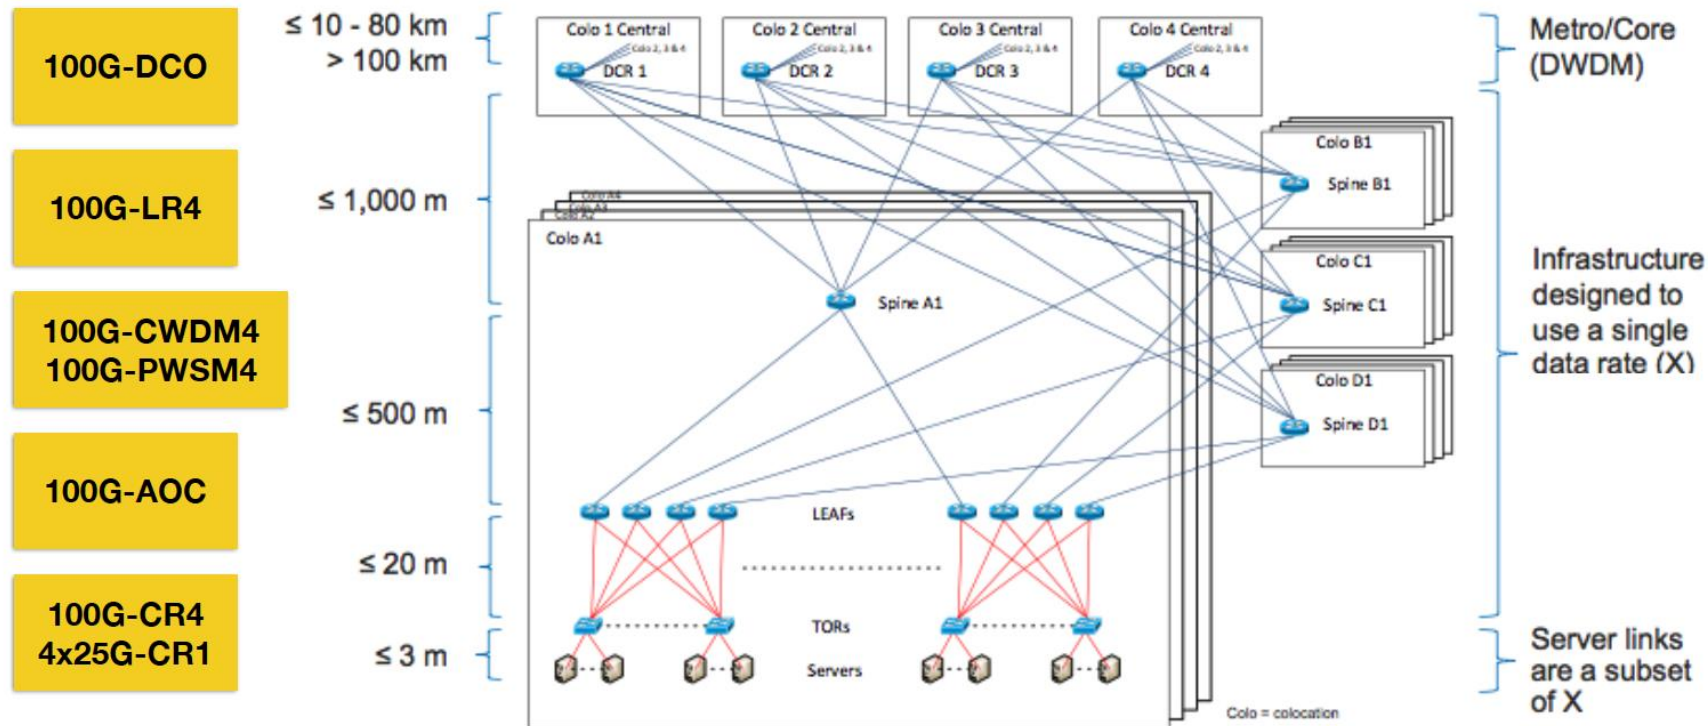
Lane Speed	10Gbps	25Gbps	50Gbps	100Gbps
1X	10G	25G	50G	100G
2X	—	50G	100G	200G
4X	40G	100G	200G	400G
8X	—	—	400G	800G
Availability	2011	2015	2018	2020

Ethernet Speed Transitions by Lane Speed


Lane Speed	10Gbps	25Gbps	50Gbps	100Gbps	
1X	10G	25G	50G	100G	Server Interface
2X	—	50G	100G	200G	
4X	40G	100G	200G	400G	Leaf-Spine Interface
8X	—	—	400G	800G	
Availability	2011	2015	2018	2020	



100G In the Cloud Network (Today)



Standard QSFP28 100GE Optics Today



- Multi-mode Fiber: QSFP-SR4
- Single-mode Fiber: QSFP-CWDM4, PSM4, LRL4 and LR4
- Cables: AOCs and Copper cables for short connections

100G ramping fast: 5 Million estimated ports in 2017

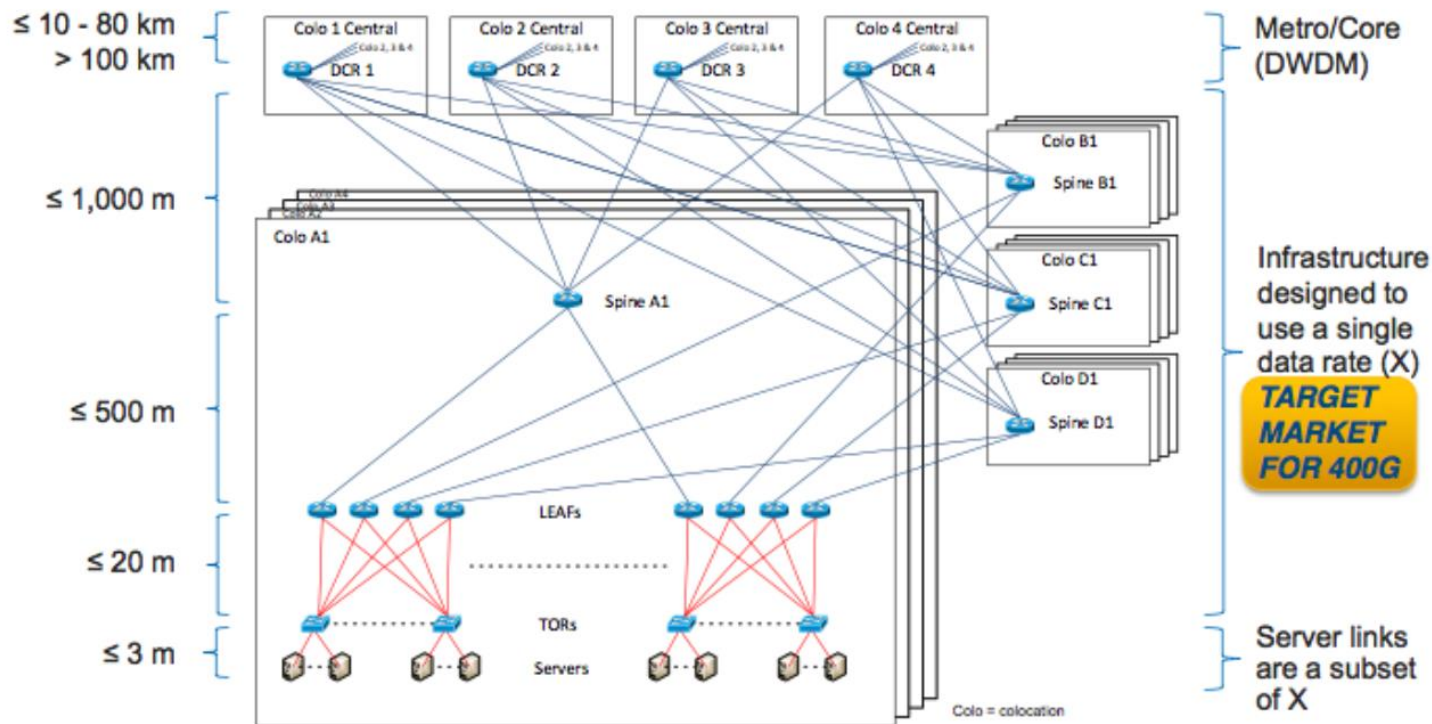
400G In the Cloud Network (starting in 2018)

400G-DCO


400G-LR8

400G-FR4
400G-DR4


400G-AOC

400G-CR8
8x50G-CR1

40G - 100G - 400G Switch Port Transition



OSFP vs CFP8 and QSFP-DD

**CFP8****QSFP-DD****OSFP**

Dimensions (mm)	CFP8	QSFP-DD	OSFP
Width	41.5	18.35	22.58
Length (overall)	107.5	89.4	107.8
Thermal Capability	12-18W	7-10W	12-15W

The OSFP (Octal Small Formfactor Pluggable)

Eight Lanes at 56 or 112Gbps

Supports 400G and 800G (2x400G)

High Port Density: 36 per 1U

14.4 Tbps with 400G, 28.8 Tbps with 800G

High Thermal Capability

Demonstrated 15W Power Capability

Accommodates full Range of Optics

Datacenter to Metro and Long Reach

Pluggable Module Form Factor

Easy to configure and service



OSFP MSA Announcement November 2016

49 Industry Leaders Collaborate to Define new 400G OSFP Transceiver Module for Next-generation Cloud Datacenter Networks

SANTA CLARA, CA. – November 15, 2016-- Arista Networks, together with Acacia Communications, Accelink, Adva Optical Networking, Amphenol, AppliedMicro, Applied Optoelectronics, Barefoot Networks, Broadcom, Cavium, ClariPhy Communications, ColorChip, Coriant, Corning, Dell EMC, Finisar, Foxconn Interconnect Technology, Fujitsu Optical Components, Google, Hewlett Packard Enterprise, Hitachi Cable Systems, Huawei Technologies, Infinera, Innolight, Innovium, Inphi, Intel, Ixia, Juniper Networks, Kaia, Lorom, Lumentum, Luxtera, Macom, Marvell, Mellanox Technologies, Molex, Multilane, NeoPhotonics, NEL America, Nokia, Oclaro, PHY-SI, SAE, Senko, Source Photonics, Sumitomo Electric Industries, TE Connectivity, and Yamaichi Electronics announced the formation of the OSFP (Octal Small Form-factor Pluggable) Multi-Source Agreement.

OSFP MSA Announcement March 2017

OSFP MSA Group Announces Form Factor Specification

Octal Small Form Factor Pluggable enables up to 800 Gbps solutions, providing an eight-fold bandwidth increase in network applications

Los Angeles – March 17, 2017 – The Octal Small Form Factor Pluggable (OSFP) Multi Source Agreement (MSA) Group has released the electrical and mechanical specifications for the new OSFP form factor. A total of 80 companies have joined the OSFP MSA to create this new standard.

Established in November 2016, the OSFP MSA group has developed a high-performance pluggable optics module form factor that is capable of supporting the full range of 400G optics technologies for datacenter and metro applications. In addition, the OSFP is designed to support the next generation of 800G optics modules that will use eight lanes of 100Gbps, and offers backwards compatibility with 100G QSFP via a simple adaptor.

OSFP 400G Use Cases



Supports all 400G use cases up to Metro and Long Reach

No single 400G optics technology addresses all market requirements

OSFP Supports Future Dual 400G and 800G Optics

Electrical and thermal performances supports eight lanes of 100G

Bulk of 400G volume will use 100G electrical interfaces

Switches with 100G lane switch silicon expected in 2020

Future High Density 100GE to Server

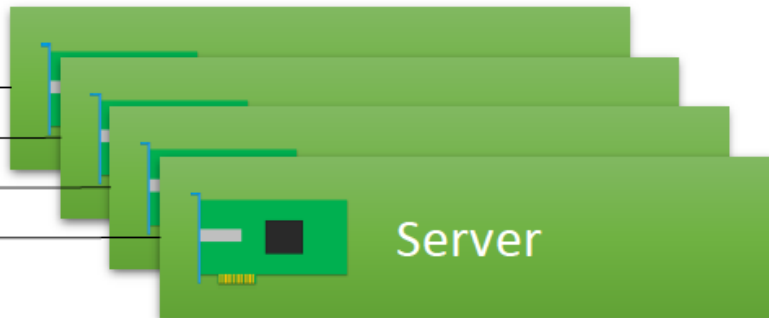


400G-DR4




MTP12 to 4x100G: SMF

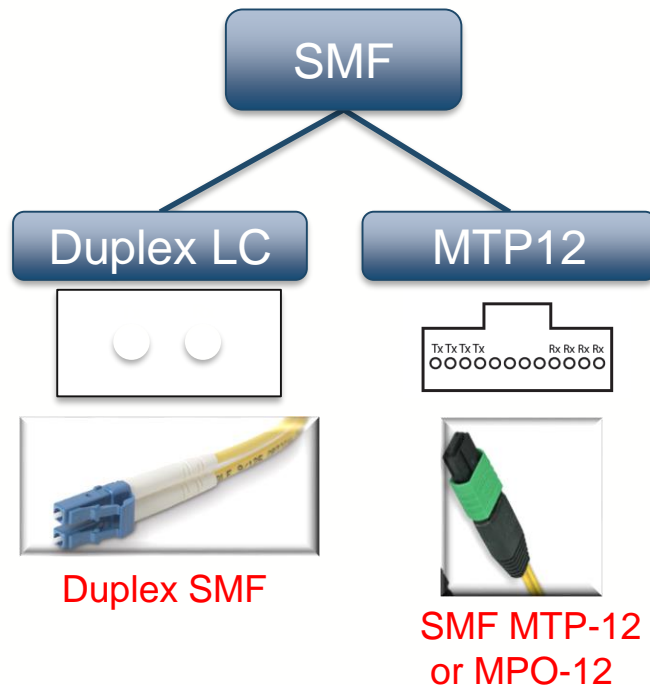
100G
100G
100G
100G



Future High Density 100GE Solution



- 400G-DR4 with 8x SMF
- 4x100G-DR1 breakout
- 500m to 2km reach
- 128 x 100GE ports per RU
- 100GE to compute/storage





400G-DCO: 100km Reach

400G-16QAM DSP + Coherent Laser

Up to 20 Terabits bandwidth per dark Fiber


Power Target < 15W




400G Coherent at the same port density as other Datacenter Optics

Other 400G/800G OSFP Transceivers

Type	Distance	Cable	Lane Speed	Power
400G-DR4	2km	8 SMF	4 x 100G	10W
400G-FR4	500m-2km	2 SMF	4 x 100G	10W
Dual 400G-DR4	2km	16 SMF	8 x 112G	12W
Dual 400G-FR4	500m-2km	4 SMF	8 x 112G	12W
800G-FR8	2km	2 SMF	8 x 112G	12W
800G-LR8	10km	2 SMF	8 x 112G	12W



Backwards Compatibility - OSFP to QSFP Adapter



OSFP Summary www.osfpmsa.org

Best Optics form factor for 400G *and* 800G (2x400G)

Electrical performance supports 112G-PAM4 lanes

Thermal performance supports modules up to 15W

Mechanical envelope supports dual 400G optics

Highest Density Pluggable Optics Transceiver

Up to 28.8 Tbps per 1U with 72 400G ports

Supports full range of Optics Technologies

Copper Cables to Metro and Long Reach DCO

Wide Industry support

80 OSFP MSA Members as of March 17, 2017



Thank You

addison@arista.com

www.arista.com